arXiv:1408.0369v7 [physics.data-an] 6 May 2015

A New Estimator of Intrinsic Dimension Based on the Multipoint Morisita Index

Jean GOLAY and Mikhail KANEVSKI

Institute of Earth Surface Dynamics, Faculty of Geosciences and Environment, University of Lausanne, 1015 Lausanne, Switzerland. Email: jean.golay@unil.ch. Phone: +41 21 692 35 41

Abstract

The size of datasets has been increasing rapidly both in terms of number of variables and number of events. As a result, the empty space phenomenon and the curse of dimensionality complicate the extraction of useful information. But, in general, data lie on non-linear manifolds of much lower dimension than that of the spaces in which they are embedded. In many pattern recognition tasks, learning these manifolds is a key issue and it requires the knowledge of their true *intrinsic dimension*. This paper introduces a new estimator of intrinsic dimension based on the multipoint Morisita index. It is applied to both synthetic and real datasets of varying complexities and comparisons with other existing estimators are carried out. The proposed estimator turns out to be fairly robust to sample size and noise, unaffected by edge effects, able to handle large datasets and computationally efficient.

Keywords: Intrinsic dimension, Multipoint Morisita index, Fractal dimension, Multifractality, Dimensionality Reduction

1. Introduction

The 21^{st} century is more and more data-dependent and, in general, when collecting data for a particular purpose, it is not known which variables matter the most. This lack of knowledge leads to the emergence of high-dimensional datasets characterized by redundant features which artificially increase the volume of data to be processed. As a result, the empty space phenomenon [1] and the curse of dimensionality [2] make it challenging to conduct pattern recognition tasks such as clustering and classification.

The goal of Dimensionality Reduction (DR) [3, 4], sometimes called manifold learning, is to address this issue by mapping the N sampled data points into the lower dimensional space where they truly lie. Such a space is often considered as a manifold of intrinsic dimension M[1] embedded in a Euclidean space of dimension E with $E \geq M$. E equals the number of variables of a dataset and the Intrinsic Dimension (ID) of a manifold is equal to the theoretical ID of the data. If a manifold is space-filling, its dimension $M \approx E$. In contrast, if the Euclidean space is partially empty, $M < E$. The optimality of DR greatly depends on the accuracy of ID estimates. An underestimation of the theoretical ID will result in the implosion of the data manifold and information will be irreparably lost. On the contrary, an overestimation will lead to noise in the final mapping. From an application perspective, DR can be used to produce low dimensional syntheses of high dimensional datasets [5] and as a preprocessing tool for supervised learning [6, 7] and data visualization [8].

DR methods perform variable transformations to capture the complex dependencies which generate redundancy within datasets. Nevertheless, it is often important not to recast data. The Fractal Dimension Reduction (FDR) algorithm [9, 10, 11, 12] was designed to this end. The fundamental idea is to remove from a dataset all the variables which do not contribute to increasing its ID. FDR can also be adapted to supervised feature selection methods [13]. The goal is then to reject irrelevant or redundant variables (or features) according to a prediction task (i.e. regression or classification). Although ID estimation lies at the core of FDR, more traditional unsupervised [14, 15, 16] and supervised [17, 18, 19, 20, 21, 22] feature selection methods do not consider it. It has, however, a great potential in speeding up search strategies, such as those used in [23, 24, 25, 26].

These different approaches highlight that ID estimation is a fundamental problem when dealing with high-dimensional datasets. Unfortunately, ID estimators [27, 28] suffer from the curse of dimensionality as well. Their overall performance depends on many factors (to various degree), such as the number of data points, the theoretical ID of data and the shape of manifolds. The present research deals with a new ID estimator in order to provide a solution to the problems raised by these factors. It is based on the recently introduced multipoint Morisita index (m-Morisita)[29, 30, 31].

[1] In Physics, mainly, M is often referred to as the degrees of freedom of data

The m-Morisita index is a measure of global clustering closely related to the concept of multifractality and, so far, it has been successfully applied within the framework of (2-dimensional) spatial data analysis [29, 30].

The paper is organized as follows: In Section 2, traditional fractal-based and maximum likelihood methods of ID estimation are presented. Section 3 derives a new ID estimator from the m-Morisita index and introduces a new algorithm for its application to high-dimensional datasets. Section 4 is devoted to comparisons between the proposed estimator and those of Section 3. Their behaviour regarding sample sizes, noise and the dimension of manifolds is analysed. A special attention is also paid to their bias and variance by using Monte-Carlo simulations and real world case studies from the UCI machine learning repository are examined. Finally, conclusions are drawn in Section 5.

2. Existing Methods

Many ID estimation methods have been proposed [27, 28, 32, 33, 34] and they can be roughly divided into projection (e.g. PCA) and geometric methods (e.g. fractal, nearest-neighbor and maximum likelihood methods). This section focuses on fractal-based and maximum likelihood estimators. They are commonly used in a wide range of applications and they generally provide non-integer values as ID estimates.

2.1. Fractal-Based Estimation Methods

The word *fractal* was first coined by B. Mandelbrot [35] to describe scale-invariant sets. At small scales δ, for a given point pattern, one has that:

$$n_{box}(\delta) \propto \delta^{-D_0} \tag{1}$$

where $n_{box}(\delta)$ is the number of grid cells necessary to cover the whole pattern and D_0 is known as the box-counting dimension [35, 36, 37]. In practical applications, due to its simplicity, D_0 often replaces the Hausdorff dimension D (or fractal dimension) and it can be proved that D_0 is an upper bound of D [38].

In complex cases, the scaling behaviour of the moments of point distributions cannot be fully characterized by only one fractal dimension and a full spectrum of generalized dimensions, D_q, is required. Such distributions

are referred to as being multifractal [39, 40, 41, 42]. D_q is generally obtained by using a generalization of the box-counting method [39, 40, 41, 43] based on Rényi's information, $RI_q(\delta)$, of q^{th} order [44]. The central scaling law of this approach can be written as follows for $q \neq 1$:

$$\exp(RI_q(\delta)) \propto \delta^{-D_q} \tag{2}$$

where

$$RI_q(\delta) = \frac{1}{1-q} \log(\sum_{i=1}^{n_{box}(\delta)} p_i(\delta)^q) \tag{3}$$

In this last equation, $p_i(\delta) = n_i/N$ is the value of the probability mass function in the i^{th} grid cell of size δ (n_i is the number of points falling into the i^{th} cell) and $q \in \mathbb{R}\backslash\{-1\}$. Finally, one has that:

$$D_q = \lim_{\delta \to 0} \frac{RI_q(\delta)}{\log(\frac{1}{\delta})} \tag{4}$$

and

$$\lim_{q \to 1} D_q = df_i \tag{5}$$

$$D_2 = df_{cor} \tag{6}$$

where df_i and df_{cor} are, respectively, the information dimension [45, 39] and the correlation dimension [46].

Usually, df_{cor} is computed with the Grassberger-Procaccia (GP) algorithm [46]. This algorithm is designed to better take advantage of the range of available pairwise distances between points. It can be introduced as follows: at small scales, for a point set, $X_N = \{x_1, \ldots, x_N\}$, one has that

$$C(\delta) \propto \delta^{df_{cor}} \tag{7}$$

where

$$C(\delta) = \frac{2}{N(N-1)} \sum_{i=1}^{N} \sum_{j=i+1}^{N} \mathbb{1}_{\{\|x_i - x_j\| \leq \delta\}} \tag{8}$$

with $\mathbb{1}$ being an indicator function and df_{cor} can be expressed as

$$df_{cor} = D_2 = \lim_{\delta \to 0} \frac{\log(C(\delta))}{\log(\delta)} \tag{9}$$

The available values of $RI_q(\delta)$ and $\log(C(\delta))$ depend on the data resolution. A commonly used method for estimating D_q consists in plotting $RI_q(\delta)$ vs $\log(\delta^{-1})$ for a chosen scale interval. The final estimate is then the slope of the linear regression fitting the linear part of the resulting chart. The procedure is the same for the GP algorithm, except that df_{cor} and $\log(C(\delta))$ replace, respectively, D_q and $RI_q(\delta)$. Eventually, both D_q (in general $0 \leqslant q \leqslant 2$) and df_{cor} can be used as ID estimators.

Although these methods may entail some disadvantages due to the finiteness of datasets [33], they have been successfully applied in various fields, such as spatial [47, 48] and time series [49] analysis, cosmology [50], climatology [36, 51, 52] and pattern recognition [53, 54]. They have also been used in different procedures improving their overall performance [55].

2.2. Maximum Likelihood Estimation Methods

The Maximum Likelihood Estimation (MLE) of ID was introduced in [28]. The proposed method relies on the assumption that the k-nearest neighbors (k-NN) of any point x_i of a point set $X_N = \{x_1, \ldots, x_N\}$ are stemming from a uniform probability density function $f(x_i)$. As a consequence, for a fixed x_i, the observations are treated as a homogeneous Poisson process within a small sphere $S_{x_i}(R)$ of radius R centred at x_i. On this basis, the inhomogeneous binomial process $\{N(t, x_i), 0 \leq t \leq R\}$ with

$$N(t, x_i) = \sum_{j=1}^{N} \mathbb{1}_{\left\{x_j \in S_{x_i}(t)\right\}} \tag{10}$$

counts the number of observations within the distance t of x_i and can be approximated as a Poisson process. The rate of this process is:

$$\lambda(t, x_i) = f(x_i)\ V(m(x_i))\ m(x_i)\ t^{m(x_i)-1} \tag{11}$$

where $m(x_i)$ is the dimension of the manifold on which x_i lies and $V(m(x_i))$ is the volume of the unit sphere in $\mathbb{R}^{m(x_i)}$ centred at x_i. The log-likelihood function of $N(t, x_i)$ can then be expressed as:

$$L(m(x_i), \theta(x_i)) = \int_0^R \log(\lambda(t, x_i))\ dN(t, x_i) - \int_0^R \lambda(t, x_i)\ dt \tag{12}$$

where $\theta(x_i) = \log(f(x_i))$. Finally, the MLE for $m(x_i)$ provides a local estimator of ID [28, 56, 57]:

$$\hat{m}_k(x_i) = \left[\frac{1}{k-2} \sum_{j=1}^{k-1} \log\left(\frac{T_k(x_i)}{T_j(x_i)}\right)\right]^{-1} \tag{13}$$

where $k > 2$ is the number of NN taken into account and $T_k(x_i)$ is the distance between x_i and its k^{th} NN.

If it is assumed that all the observations belong to the same manifold, one has that:

$$\hat{m}_k = \frac{1}{N} \sum_{i=1}^{N} \hat{m}_k(x_i) \tag{14}$$

which is simply an average over the whole dataset and, for $k \in \{k_1, k_1 + 1, ..., k_2\}$ with $k_1 > 2$, the final estimate of ID is provided in [28]:

$$\hat{m} = \frac{1}{k_2 - k_1 + 1} \sum_{k=k_1}^{k_2} \hat{m}_k \tag{15}$$

A delicate issue which arises from Equations 13, 14 and 15 is the range of the values k to be chosen. In practical applications, this is similar to the choice of the scale interval in the fractal-based methods. Here, as well, the finiteness of datasets may greatly influence the final estimate of ID if the considered values k are not carefully selected. In [28], it is advocated to retain a range of small to moderate values, so that each $S_{x_i}(R)$ is small enough to ensure $f(x_i) \approx const$ and large enough to contain sufficiently many points.

In [58], a modified version of the MLE algorithm is proposed. It consists in averaging the inverse of the N estimators $\hat{m}_k(x_i)$ of Equation 14, so that the final estimator of Equation 15 is replaced with:

$$\hat{m} = \frac{1}{k_2 - k_1 + 1} \sum_{k=k_1}^{k_2} \left[\frac{1}{N(k-1)} \sum_{i=1}^{N} \sum_{j=1}^{k-1} \log \left(\frac{T_k(x_i)}{T_j(x_i)} \right) \right]^{-1} \tag{16}$$

Although the second version of the algorithm is better for small values of k, both of them yield similar results [57] and have been successfully applied in various studies [28, 58, 56]. Finally, notice that, in the remainder of this paper, the estimators of Equations 15 and 16 will be named after their authors, $\hat{m}_{LB}$ and $\hat{m}_{MG}$.

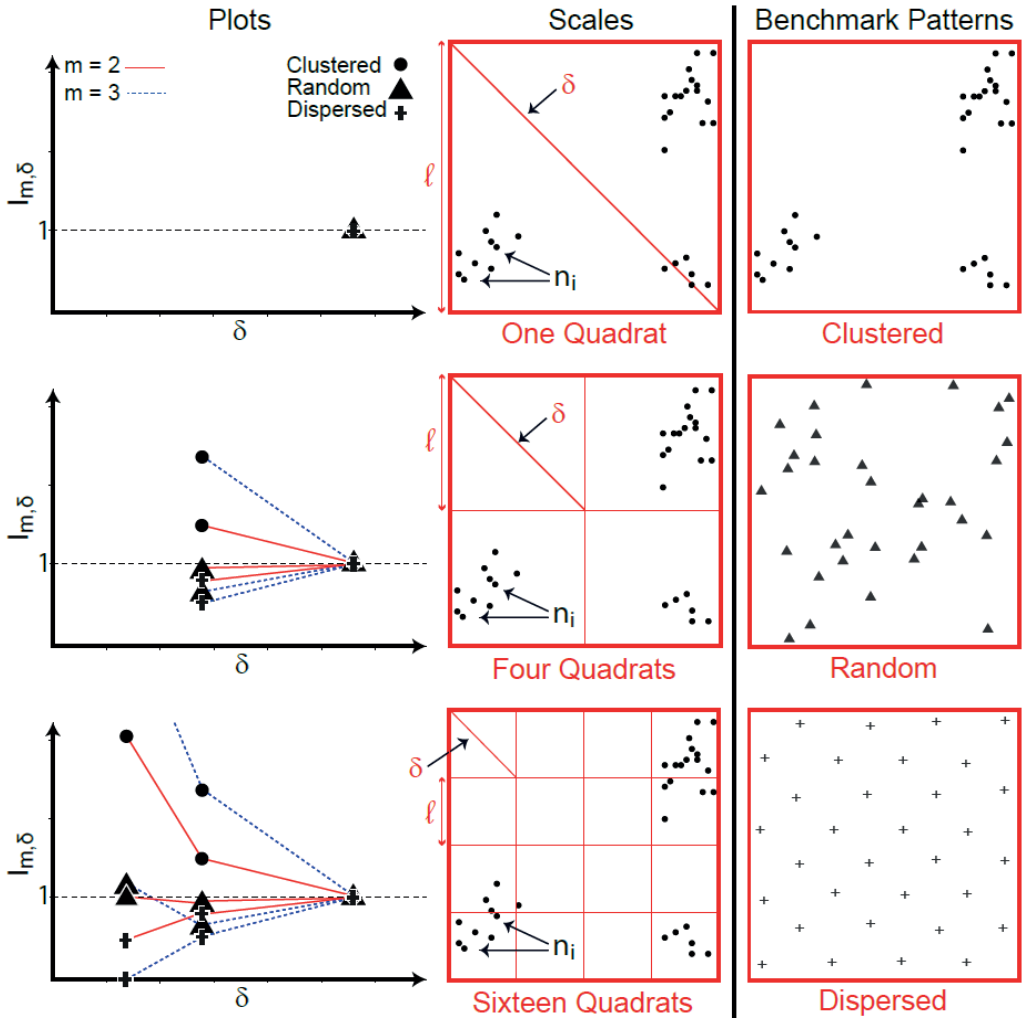

Figure 1: Computation of the m-Morisita index in two dimensions for $m = 2$ and $m = 3$ and for three benchmark patterns

3. A New Estimator of Intrinsic Dimension

3.1. The m-Morisita Index

The m-Morisita index [29, 30], $I_{m,\delta}$, is a global measure of clustering. It is a generalization of the Morisita index [31, 59] and it was first proposed in [30] for the analysis of population distributions in ecology. It was later modified in [29] to take into account the notion of scale in spatial data anal-

ysis and a relationship to multifractality was established.

For its computation, $I_{m,\delta}$, requires a dataset to be covered with a grid of Q quadrats (or cells) of changing size δ (see Figure 1). For a fixed δ, $I_{m,\delta}$ measures how many times more likely it is that m ($m \geq 2$) randomly selected data points will be from the same quadrat than it would be if the N points of the dataset were distributed at random (see Figure 1). It is calculated as follows:

$$I_{m,\delta} = Q^{m-1} \frac{\sum_{i=1}^{Q} n_i(n_i-1)(n_i-2)\cdots(n_i-m+1)}{N(N-1)(N-2)\cdots(N-m+1)} \tag{17}$$

where n_i is the number of points in the i^{th} quadrat and N is the total number of points. The computation of the index starts with a relatively large quadrat size δ. It is then reduced until it reaches a minimum value and a plot relating every $I_{m,\delta}$ to its matching δ can be drawn.

Figure 1 illustrates the computation of the index in two dimensions for three benchmark point distributions (or patterns), for three different scales and for $m = 2$ and $m = 3$. For the highest possible δ, when only one quadrat is considered, $I_{m,\delta}$ returns the same value for each pattern and m. As the number of quadrats increases, $I_{m,\delta}$ adopts a specific behaviour for each of the three benchmark distributions. If the points are distributed at random, every computed $I_{m,\delta}$ oscillates around the value of 1. If the points are clustered, the value of the index increases as δ decreases and, finally, if the points are dispersed, the index approaches 0 at small scales [29, 60]. Further, as m increases, $I_{m,\delta}$ becomes more and more sensitive to the structure of the pattern under study. In complex situations, $I_{m,\delta}$ computed with small m may miss structures which are detected with higher m.

3.2. The Morisita Estimator of Intrinsic Dimension

Several parallels can be drawn between the m-Morisita index and Rényi's information of q^{th} order. In particular, it was established, for fractal point sets, that [29]:

$$\lim_{\delta \to 0} \frac{\log\left(I_{m,\delta}\right)}{\log(\frac{1}{\delta})} \frac{1}{m-1} \approx E - D_m = C_m \tag{18}$$

where $m \in \{2, 3, 4, \cdots\}$, C_m is the codimension of order $q = m$, E is the dimension of the Euclidean space where the dataset is embedded and D_m is

Rényi's generalized dimension of order $q = m$. In practical applications, for finite datasets, it can be shown that Equation 18 is verified only under the condition that $H := \underset{i}{max}\,(n_i) \gg m$ [29]. If so, C_m can be estimated from the slope, S_m, of the straight line fitting the linear part of the plot relating $\log\,(I_{m,\delta})$ to $\log\left(\frac{1}{\delta}\right)$. Then, $C_m \approx S_m/(m-1)$ and one has that:

$$D_m \approx E - \left(\frac{S_m}{m-1}\right) \tag{19}$$

where $m \in \{2, 3, 4, \cdots\}$.

In high-dimensional spaces, the condition $H \gg m$ is hardly ever met at small scales. In such situations, it is important to notice that the major difference between $\log\,(I_{m,\delta})$ and $RI_m(\delta)$ lies in the following inequality:

$$n_i^m > n_i(n_i-1)\cdots(n_i-m+1) \tag{20}$$

Consequently, unlike $\log\,(I_{m,\delta})$, $RI_m(\delta)$ seriously overestimates the probability of randomly drawing m-tuples of points from grid cells characterized by a small n_i. Such cells are numerous at small scales or when the sample size is limited and can greatly affect the accuracy of D_m. From this perspective, it is possible to suggest a new ID estimator based on $\log\,(I_{m,\delta})$ (for $m \geq 2$):

$$M_m := E - \left(\frac{S_m}{m-1}\right) \tag{21}$$

which should be more robust to sample size than D_m (see Equation 4). Notice that M_m should only be computed under the condition that $H > m$ at all considered scales. In the remainder of this paper, M_m will be referred to as the Morisita estimator of ID. It will be thoroughly tested in Section 4 with synthetic datasets of various complexities and with real data from the UCI machine learning repository.

3.3. An Algorithm for Large Datasets

Algorithms used for handling large datasets should be affected as little as possible by the amount of main memory available. Regarding M_m, the main issue concerns the way the number of points per quadrat is counted. A good algorithm must be able to effectively disregard empty quadrats. It is also more appealing if its implementation is straightforward in most programming environments (e.g. R and Matlab). To fulfil these requisites, the

Algorithm 1 The Morisita INDex for ID estimation (MINDID)

INPUT: a $N \times E$ matrix, D, with N points and E features; a vector L of values ℓ; a vector M of contiguous values m. **OUTPUT:** a vector containing M_m for each value m.

1: Rescale each variable to $[0, 1]$
2: **for all** values ℓ **do**
3: Divide each element of D by ℓ and round the result to the next lowest whole number in a $N \times E$ matrix called $D\ell$
4: Count the number, nbr_l, of different lines in $D\ell$ and store their frequency in a vector, ni, of size nbr_l
5: **for all** values m **do**
6: Compute $\log(I_{m,\delta})$ (See Equation 23) using the vector ni and store the result in a $\mid L \mid \times \mid M \mid$ matrix called logMindex
7: **end for**
8: **end for**
9: Optional: Compute the values δ (See Equation 24) using the vector L and store the result in a vector Δ of cardinality $\mid L \mid$
10: **for all** values m **do**
11: Compute M_m (See Equation 21) using logMindex and Δ (or L)
12: **end for**

MINDID algorithm (see Algorithm 1) incorporates a version of an algorithm suggested in [10]. It rescales each variable of a dataset D to $[0, 1]$ and takes advantage of the properties of the square cells of the hyper-grid covering the data: for a given cell size δ, each value of the dataset is divided by ℓ (see Figure 1) and rounded to the next lowest integer. In this way, in the resulting matrix, all the data points falling into the same cell are matched by as many equal lines.

Another issue concerns Q^{m-1}, since it is often given the value *Inf* (e.g. R and Matlab) for small values of ℓ when $E \gg 1$. A way to overcome this problem is to resort to $\log(I_{m,\delta})$ instead of $I_{m,\delta}$. Q^{m-1} is related to ℓ and E through:

$$Q^{m-1} = \left(\frac{1}{\ell}\right)^{E(m-1)} \tag{22}$$

and

$$\log(I_{m,\delta}) = E(1-m)\log(\ell) + \log\left(\frac{\sum_{i=1}^{Q} n_i(n_i-1)\cdots(n_i-m+1)}{N(N-1)\cdots(N-m+1)}\right) \tag{23}$$

This solution is satisfactory, since the computation of M_m (See Equation 21) only requires $\log(I_{m,\delta})$. Consequently, the second part of the MINDID algorithm is devoted to the implementation of Equation 23.

Finally, the computation of M_m is carried out using either $\log(\ell)$ or $log(\delta)$. In the rest of this paper, the second option will be preferred. Notice that ℓ and δ are related as follows:

$$\delta = \ell\sqrt{E} \tag{24}$$

4. Assessment of the Morisita Estimator of ID

4.1. Synthetic Data

Several datasets were built, so that each of them resides on a known manifold (or near a known manifold in the case of noisy data). They can be divided into four categories (see Figure 2):

1. Swiss rolls (e.g. [3]) of 1000, 5000 and 10 000 points. The theoretical ID of the data is equal to 2.
2. Noisy Swiss rolls of 1000, 5000 and 10 000 points. The noise is modelled as a Gaussian variable $G \sim N(0, \sigma^2)$ where σ varies from 0 to 0.5.
3. Uniform clouds of 1000, 5000 and 10 000 points. Each of the N distinct points x_i, with $i \in \{1, 2, \cdots, N\}$, is described by a E-dimensional vector $[x_i^1, x_i^2, \cdots, x_i^E]^T \in \mathbb{R}^E$, the components of which are sampled from E i.i.d. variables following a uniform distribution. E is gradually increased from 1 to 7 and is equal to the theoretical ID of the data.
4. This last category is based on the properties of the Cartesian product of some fractals [38, 61]. One-dimensional Cantor sets of 8192 and 65 536 points are first created. The resulting vectors are shuffled seven times to generate as many variables. A Euclidean space $\mathbb{R}^E$ can then be constructed and E is gradually increased from 1 to 7. The dimension of the data manifold (i.e. the theoretical ID) is equal to $\frac{\log(2)}{\log(3)}E$, where $\frac{\log(2)}{\log(3)}$ is the Hausdorff dimension of a one-dimensional Cantor set.

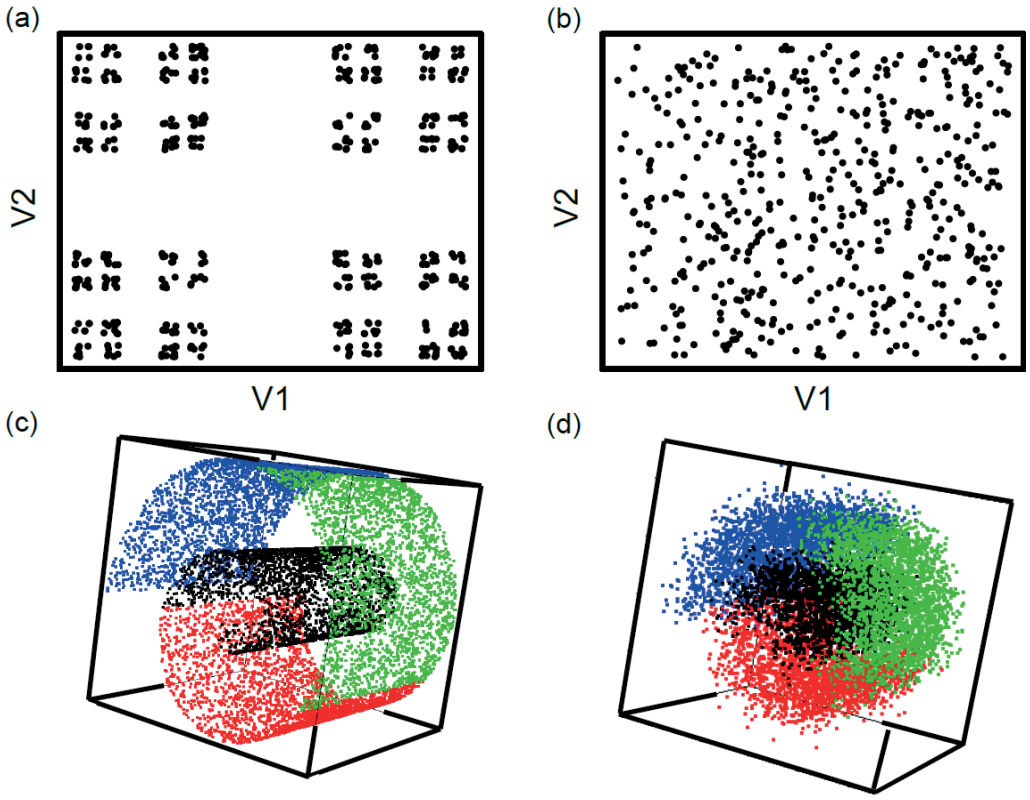

Figure 2: (a) A Cantor set of 512 points for $E = 2$; (b) a realization of a uniform cloud of 500 points for $E = 2$; (c) a Swiss roll of 10 000 points without noise and (d) a Swiss roll of 10 000 points with noise $G \sim N(0, 0.25^2)$

4.2. Comparison Between M_m and D_m Using Synthetic Data

Two categories of datasets were used to compare M_m with D_m: the uniform point clouds and the Cantor sets. They were employed with varying N and E (see Subsection 4.1) and for each combination of these two parameters, 100 sets were generated. The results are displayed in Figures 3 and 4 and in Tables 1 and 2.

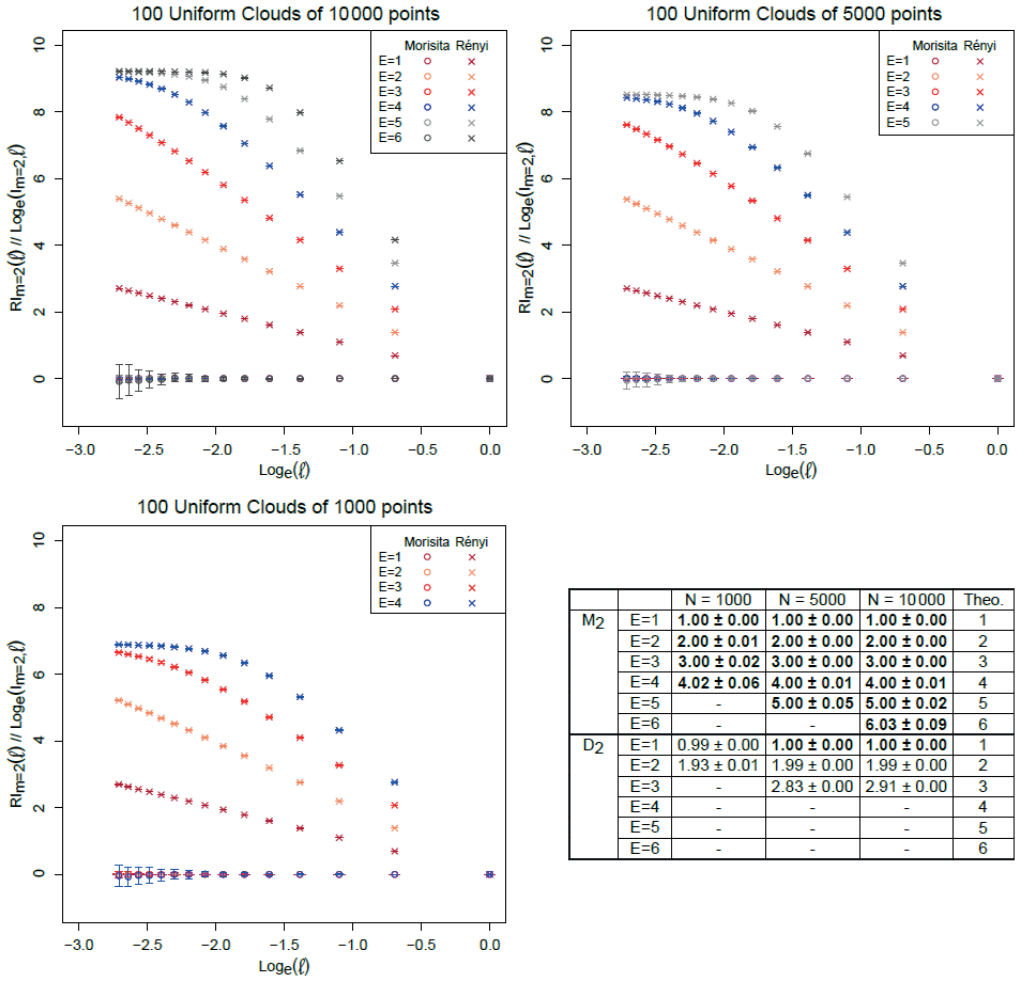

		N = 1000	N = 5000	N = 10000	Theo.
M_2	E=1	**1.00 ± 0.00**	**1.00 ± 0.00**	**1.00 ± 0.00**	1
	E=2	**2.00 ± 0.01**	**2.00 ± 0.00**	**2.00 ± 0.00**	2
	E=3	**3.00 ± 0.02**	**3.00 ± 0.00**	**3.00 ± 0.00**	3
	E=4	**4.02 ± 0.06**	**4.00 ± 0.01**	**4.00 ± 0.01**	4
	E=5	-	**5.00 ± 0.05**	**5.00 ± 0.02**	5
	E=6	-	-	**6.03 ± 0.09**	6
D_2	E=1	0.99 ± 0.00	**1.00 ± 0.00**	**1.00 ± 0.00**	1
	E=2	1.93 ± 0.01	1.99 ± 0.00	1.99 ± 0.00	2
	E=3	-	2.83 ± 0.00	2.91 ± 0.00	3
	E=4	-	-	-	4
	E=5	-	-	-	5
	E=6	-	-	-	6

Figure 3: The results of the application of $I_{2,\ell}$ and $RI_2(\ell)$ to 100 uniform point clouds of 10 000, 5000 and 1000 points. In the bottom-right table, the corresponding estimates of D_2 and M_2 are provided as follows: the mean (computed over the 100 sets) $\pm$ the standard deviation. The results written in bold script indicate that the theoretical ID falls within the mean value $\pm$ one standard deviation.

In the case of the uniform point clouds, $RI_2(\ell)$ and $\log\left(I_{2,\ell}\right)$ were computed by using an interval of the parameter ℓ^{-1} ranging from 1 to 15 (see Equation 24). Figure 3 shows the results for $N = 10\,000$, $N = 5000$ and $N = 1000$ and for increasing E. The points of each plot are the mean values yielded by the two indices over the 100 sets and the error bars correspond to the standard deviations.

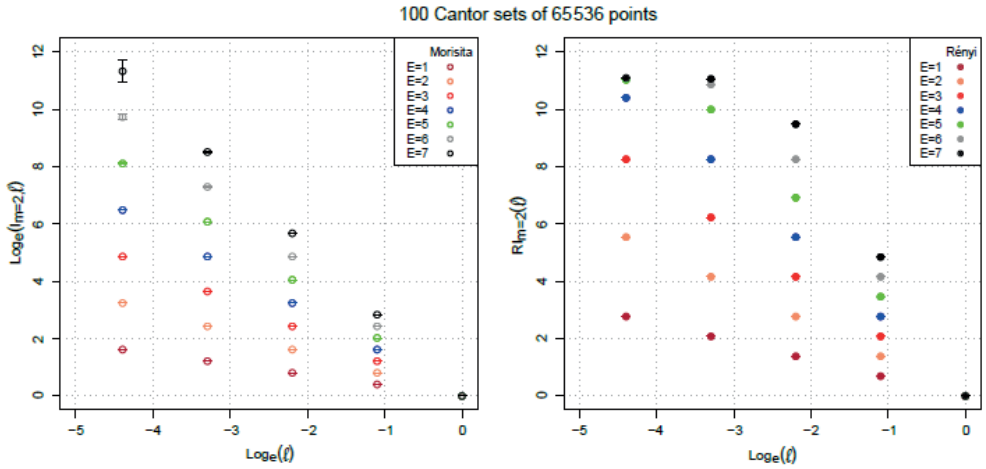

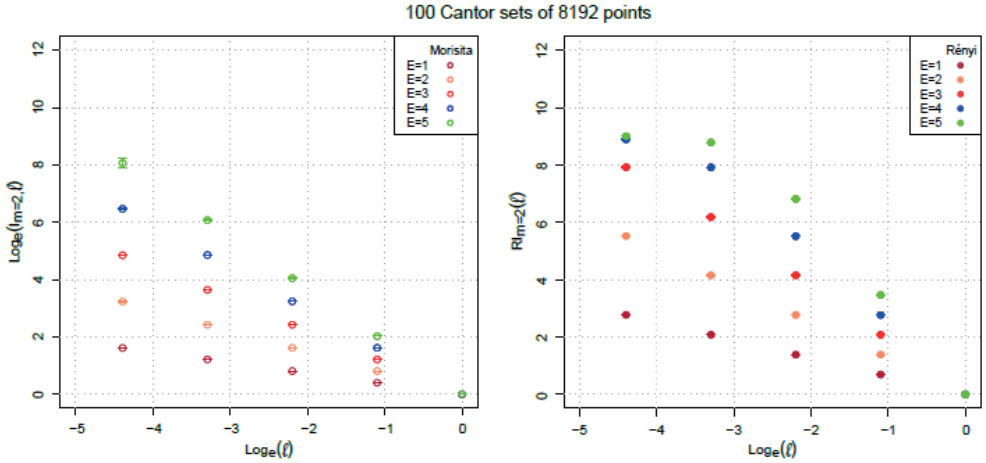

		N = 8192	N = 65536	Theo.
M_2	E=1	**0.63 ± 0.00**	**0.63 ± 0.00**	0.63
	E=2	**1.26 ± 0.00**	**1.26 ± 0.00**	1.26
	E=3	1.90 ± 0.00	**1.89 ± 0.00**	1.89
	E=4	**2.53 ± 0.01**	**2.52 ± 0.00**	2.52
	E=5	**3.16 ± 0.04**	3.16 ± 0.00	3.15
	E=6	-	**3.79 ± 0.02**	3.79
	E=7	-	**4.43 ± 0.07**	4.42
D_2	E=1	**0.63 ± 0.00**	**0.63 ± 0.00**	0.63
	E=2	**1.26 ± 0.00**	**1.26 ± 0.00**	1.26
	E=3	1.82 ± 0.00	1.88 ± 0.00	1.89
	E=4	-	2.39 ± 0.00	2.52
	E=5	-	-	3.15
	E=6	-	-	3.79
	E=7	-	-	4.42

Figure 4: The results of the application of $I_{2,\ell}$ (left) and $RI_2(\ell)$ (right) to 100 Cantor sets of 8192 and 65 536 points. In the bottom table, the corresponding estimates of D_2 and M_2 are provided as follows: the mean (computed over the 100 sets) ± the standard deviation. The results written in bold script indicate that the theoretical ID falls within the mean value ± one standard deviation.

If the value of H (see Subsection 3.2) was smaller than 2 at certain scales ℓ for a given dimension E, the entire plots describing the corresponding behaviours of $RI_2(\ell)$ and $\log(I_{2,\ell})$ were not drawn. This was motivated by the condition that H must be greater than m at all scales for M_m to be computed.

In spite of this limitation, M_m provides better ID estimates than D_q when the same values of ℓ are considered. In Figure 3, this is highlighted, for the largest E, by the steady state reached by $RI_2(\ell)$ at small scales, which shows a departure from the power law of Equation 2. As a consequence, D_2 cannot be derived from a linear regression calculated over the whole range of ℓ. In contrast, $\log(I_{2,\ell})$ follows, on average and throughout the scales, the power law which underlies Equation 21: all the plots are superimposed on a constant mean level of 0 as expected from a Poisson distribution (i.e. a random space-filling set). The cost of this near absence of bias is an increase in the variability of the values provided by $\log(I_{2,\ell})$ as ℓ decreases. As indicated in the table of Figure 3, this drawback has only a small impact on the variabilities of the final ID estimates which remain low.

Still with regard to the table of Figure 3, the means and standard deviations of the ID estimates[2] were calculated only if the dependence between $\log(I_{2,\ell})$ or $RI_2(\ell)$ and ℓ could be reasonably approximated by using a linear regression over all the scales (i.e if only one slope could be distinguished in the different plots). The results show that D_2 becomes unreliable for $E > 2$, while M_2 works better and can even be used up to $E = 6$ for $N = 10\,000$.

The results of M_2 and D_2 were written in bold script if the theoretical ID fell within the mean value $\pm$ one standard deviation. Shapiro-Wilk tests were first conducted to check whether or not the estimates could be assumed to come from normal distributions (only if the standard deviations were not equal to 0.00). At a 5% α level, the hypothesis of normality could not be rejected for any of the estimate distributions. The theoretical ID fell within the mean value $\pm$ one standard deviation for each result yielded by M_2. Regarding D_2, it only happened twice for high values of N and for $E = 1$. Notice that the same notation will be used throughout the section.

[2]The linear slope of the plots must be multiplied by -1 to yield S_m and D_q, since the x-axis represents ℓ instead of ℓ^{-1}

$N = 65\,536$	M_2	D_2	M_3	D_3	M_5	D_5
$E = 1$	$\mathbf{0.63 \pm 0.00}$	$\mathbf{0.63 \pm 0.00}$	$\mathbf{0.63 \pm 0.00}$	$\mathbf{0.63 \pm 0.00}$	$\mathbf{0.63 \pm 0.00}$	$\mathbf{0.63 \pm 0.00}$
$E = 2$	$\mathbf{1.26 \pm 0.00}$	$\mathbf{1.26 \pm 0.00}$	$\mathbf{1.26 \pm 0.00}$	1.26 ± 0.00	$\mathbf{1.26 \pm 0.00}$	$\mathbf{1.26 \pm 0.00}$
$E = 3$	$\mathbf{1.89 \pm 0.00}$	1.88 ± 0.00	$\mathbf{1.89 \pm 0.00}$	1.88 ± 0.00	$\mathbf{1.89 \pm 0.00}$	1.87 ± 0.00
$E = 4$	$\mathbf{2.52 \pm 0.00}$	2.39 ± 0.00	$\mathbf{2.52 \pm 0.00}$	2.37 ± 0.00	$\mathbf{2.53 \pm 0.01}$	2.33 ± 0.00
$E = 5$	3.16 ± 0.00	-	$\mathbf{3.16 \pm 0.02}$	-	-	-
$E = 6$	$\mathbf{3.79 \pm 0.02}$	-	-	-	-	-
$E = 7$	$\mathbf{4.43 \pm 0.07}$	-	-	-	-	-

Table 1: The results of the application of M_m and D_m to 100 Cantor sets of 65 536 points for $m = 2$, $m = 3$ and $m = 5$. The mean $\pm$ the standard deviation is provided for each E-dimensional space. The results written in bold script indicate that the theoretical ID falls within the mean value $\pm$ one standard deviation.

Finally, the number of E-dimensional spaces, for which M_2 can be calculated, decreases as N is reduced. The reason is that it becomes less likely that at least two points will fall into the same cell at small scales when $E \gg 1$. Nevertheless, whatever N, the bias affecting $RI_2(\ell)$ is always noticeable for the greatest E and tends to lead to an underestimation of ID in any case.

Similar comments can be made about the results obtained for the Cantor sets (see Figure 4). For this second category of data, the interval of the parameter ℓ^{-1} follows a geometric series with ratio $r = 3$ and ranges from 1 to 81. In this way, the grid used for the computation of both $\log(I_{2,\ell})$ and $RI_2(\ell)$ is in accordance with the mathematical construction of a Cantor set. For the same arguments as those previously set out, M_2 turns out to be a more reliable estimator of ID than D_2: here as well, a bias affects the behaviour of $RI_2(\ell)$ when the number of points in occupied cells is low. It is also interesting to notice that M_2 can be computed up to $E = 5$ for $N = 8192$, although the lowest considered scale ℓ is smaller than it was for the uniform point clouds. This is due to the dimension of the data manifold that is systematically less than E. As a consequence, a Cantor set is not space-filling and it is more likely that at least two points will fall into the same cell than it would be if the N data points were randomly distributed within the entire Euclidean space. This observation highlights that the curse of ID [12] (i.e. the problems induced by a high ID) is a central issue when studying high-dimensional spaces.

The Cantor sets were also used to assess the accuracy of M_m and D_m for m greater than 2. Tables 1 and 2 show the results for, respectively, $N = 65\,536$ and $N = 8192$. In both cases, M_m provides results closer

$N = 8192$	M_2	D_2	M_3	D_3	M_5	D_5
$E = 1$	$\mathbf{0.63 \pm 0.00}$	$\mathbf{0.63 \pm 0.00}$	$\mathbf{0.63 \pm 0.00}$	$\mathbf{0.63 \pm 0.00}$	$\mathbf{0.63 \pm 0.00}$	$\mathbf{0.63 \pm 0.00}$
$E = 2$	$\mathbf{1.26 \pm 0.00}$	$\mathbf{1.26 \pm 0.00}$	$\mathbf{1.26 \pm 0.00}$	1.25 ± 0.00	$\mathbf{1.26 \pm 0.00}$	1.25 ± 0.00
$E = 3$	1.90 ± 0.00	1.82 ± 0.00	1.90 ± 0.00	1.80 ± 0.00	1.90 ± 0.01	1.76 ± 0.00
$E = 4$	$\mathbf{2.53 \pm 0.01}$	-	$\mathbf{2.53 \pm 0.03}$	-	-	-
$E = 5$	$\mathbf{3.16 \pm 0.04}$	-	-	-	-	-
$E = 6$	-	-	-	-	-	-
$E = 7$	-	-	-	-	-	-

Table 2: The results of the application of M_m and D_m to 100 Cantor sets of 8192 points for $m = 2$, $m = 3$ and $m = 5$. The mean $\pm$ the standard deviation is provided for each E-dimensional space. The results written in bold script indicate that the theoretical ID falls within the mean value $\pm$ one standard deviation.

to the theoretical ID than those yielded by D_m for $E > 2$. Nevertheless, the difference between the two estimators is lessened as the value of m is increased. This follows from Equation 20: when m ($m = q$) is high, the impact of the inequality is negligible, while the implementation of M_m requires more data points than that of D_m.

4.3. Comparison Between M_2 and the Distance-Based Estimators Using Synthetic Data

M_2 was also compared to the other estimators of ID presented in Section 2, namely d_{cor}, $\hat{m}_{LB}$ and $\hat{m}_{MG}$. The Swiss rolls and the uniform point clouds were used for this task. The parameters of each estimator were set using the Swiss rolls of 1000 points and stayed unchanged in the whole subsection. At each step, it was made sure that these parameters were close to the ideal ones. Regarding M_2, the interval of the parameter ℓ^{-1} was chosen, so that it ranged from 5 to 15 and the two MLE estimators were computed with k going from 10 to 20. More challenging, the interval of the parameter δ of d_{cor} turned out to be relatively complicated to set, since it tended to deviate from the ideal values as both N and E were increased. It was finally decided to resort to percentiles and the problem was empirically solved as follows: (a) 1 percent of the pairwise distances between points had to be lower than the smallest δ and 7 percent of them had to be lower than the largest one; (b) the range of the interval was divided by 100 to produce intermediate values of δ. It is also worth mentioning that each variable was rescaled, so that it ranged from 0 to 1. Such a transformation is mandatory when working with data of different nature. Regarding d_{cor}, $\hat{m}_{LB}$ and $\hat{m}_{MG}$, it amounts to using the Mahalanobis distance and, if the data are not noisy, it has, of course, no influence on M_m (or D_q), since the grid employed in its computation is also transformed.

Swiss Roll	$N = 1000$	$N = 5000$	$N = 10\,000$
M_2	$\mathbf{2.03 \pm 0.03}$	2.03 ± 0.01	2.03 ± 0.00
d_{cor}	1.95 ± 0.01	1.94 ± 0.00	1.94 ± 0.00
$\hat{m}_{LB}$	1.95 ± 0.02	1.98 ± 0.01	1.98 ± 0.01
$\hat{m}_{MG}$	1.95 ± 0.02	1.98 ± 0.01	1.98 ± 0.00

Table 3: The results of the application of the four estimators to Swiss rolls of $N = 1000$, $N = 5000$ and $N = 10\,000$ points (100 sets for each N). The mean $\pm$ the standard deviation of the estimators, computed over the 100 sets, is provided for each N. The result written in bold script indicates that the theoretical ID falls within the mean value $\pm$ one standard deviation (the theoretical ID of the Swiss Roll is 2).

Tables 3 and 4 show the results (mean $\pm$ standard deviation computed over 100 sets) provided by the four estimators for, respectively, the Swiss rolls and the uniform point clouds. The values yielded by the four estimators are similar when the data points are densely distributed on their manifold. This situation is encountered for the Swiss rolls and for relatively low E and high N in the case of the uniform distributions. When the data points are sparse (low N and/or high theoretical ID) and provided that it can be calculated, M_2 provides better ID estimates. d_{cor}, $\hat{m}_{LB}$ and $\hat{m}_{MG}$ are always able to yield a result, but, when the data are simply too sparse for M_2 to be computed, they tend to seriously underestimate the true ID of the data. The underestimation reduces as N increases and it can be assumed that it is caused (at least partially) by edge effects.

Several Edge effect corrections have been proposed and thoroughly studied in spatial data analysis [62], but the problem has often been overlooked in ID estimation methods. d_{cor} appears to be more affected by this problem than the other estimators. Nevertheless, the case of d_{cor} is difficult to deal with, since the ideal range of the parameter δ is very sensitive to N and E. Consequently, concerning the uniform distribution clouds, it was decided that a new series of 10 estimations would be carried out for $N = 5000$ and $N = 10\,000$ with improved parameters. The results appear in italics in Table 4 and reveal that the modifications have improved the estimates, even if $\hat{m}_{LB}$ and $\hat{m}_{MG}$ stay better.

		M_2	d_{cor}	$\hat{m}_{LB}$	$\hat{m}_{MG}$
$N = 1000$	$E = 1$	$\mathbf{1.00 \pm 0.00}$	$\mathbf{0.99 \pm 0.01}$	$\mathbf{1.00 \pm 0.01}$	$\mathbf{1.00 \pm 0.01}$
	$E = 2$	$\mathbf{2.00 \pm 0.01}$	1.91 ± 0.01	1.95 ± 0.02	1.95 ± 0.02
	$E = 3$	$\mathbf{3.00 \pm 0.05}$	2.74 ± 0.02	2.83 ± 0.02	2.83 ± 0.03
	$E = 4$	$\mathbf{4.04 \pm 0.19}$	3.51 ± 0.03	3.66 ± 0.03	3.65 ± 0.04
	$E = 5$	-	4.24 ± 0.04	4.45 ± 0.05	4.44 ± 0.05
	$E = 6$	-	4.93 ± 0.04	5.20 ± 0.06	5.20 ± 0.06
	$E = 7$	-	5.59 ± 0.04	5.92 ± 0.07	5.92 ± 0.07
$N = 5000$	$E = 1$	$\mathbf{1.00 \pm 0.00}$	0.99 ± 0.00 $\boldsymbol{1.00 \pm 0.00}$	$\mathbf{1.00 \pm 0.00}$	$\mathbf{1.00 \pm 0.00}$
	$E = 2$	$\mathbf{2.00 \pm 0.00}$	1.91 ± 0.01 1.95 ± 0.00	1.98 ± 0.01	1.98 ± 0.01
	$E = 3$	$\mathbf{3.00 \pm 0.01}$	2.74 ± 0.01 2.84 ± 0.01	2.90 ± 0.01	2.90 ± 0.01
	$E = 4$	$\mathbf{4.00 \pm 0.04}$	3.51 ± 0.01 3.68 ± 0.01	3.78 ± 0.02	3.77 ± 0.02
	$E = 5$	$\mathbf{5.02 \pm 0.13}$	4.24 ± 0.01 4.47 ± 0.02	4.61 ± 0.02	4.60 ± 0.02
	$E = 6$	-	4.93 ± 0.01 5.22 ± 0.01	5.40 ± 0.03	5.40 ± 0.03
	$E = 7$	-	5.58 ± 0.01 5.94 ± 0.01	6.18 ± 0.03	6.17 ± 0.03
$N = 10\,000$	$E = 1$	$\mathbf{1.00 \pm 0.00}$	0.99 ± 0.00 $\boldsymbol{1.00 \pm 0.00}$	$\mathbf{1.00 \pm 0.00}$	$\mathbf{1.00 \pm 0.00}$
	$E = 2$	$\mathbf{2.00 \pm 0.00}$	1.91 ± 0.00 1.97 ± 0.00	1.98 ± 0.00	1.98 ± 0.01
	$E = 3$	$\mathbf{3.00 \pm 0.00}$	2.74 ± 0.01 2.90 ± 0.01	2.93 ± 0.01	2.92 ± 0.01
	$E = 4$	$\mathbf{4.00 \pm 0.02}$	3.51 ± 0.01 3.78 ± 0.01	3.82 ± 0.01	3.81 ± 0.01
	$E = 5$	$\mathbf{5.00 \pm 0.08}$	4.24 ± 0.01 4.62 ± 0.01	4.66 ± 0.02	4.66 ± 0.02
	$E = 6$	$\mathbf{6.06 \pm 0.24}$	4.93 ± 0.01 5.42 ± 0.01	5.48 ± 0.02	5.47 ± 0.02
	$E = 7$	-	5.58 ± 0.01 6.19 ± 0.02	6.27 ± 0.02	6.26 ± 0.02

Table 4: The results of the application of the four estimators to uniform distribution clouds of $N = 1000$, $N = 5000$ and $N = 10\,000$ points (100 sets for each N). The mean $\pm$ standard deviation, computed over the 100 sets, is provided for each N. The results written in bold script indicate that the theoretical ID falls within the mean value $\pm$ one standard deviation (the theoretical ID of a uniform point cloud is E).

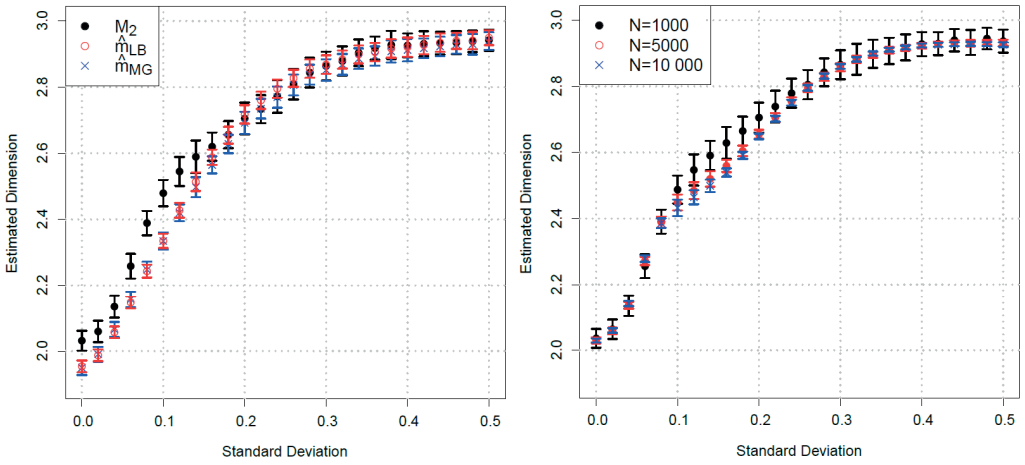

Figure 5: (left) Application of M_2, $\hat{m}_{LB}$ and $\hat{m}_{MG}$ to noisy Swiss rolls of $N = 1000$; (right) Application of M_2 to noisy Swiss rolls of $N = 1000$, $N = 5000$ and $N = 10\,000$. The x-axis represents the standard deviation of the noise and, for each level, 100 sets are considered.

4.4. Comparison Between the Estimators Using Noisy Synthetic Data

The presence of noise implies that the data points are located near a manifold instead of being exactly on it [28]. Consequently, a robust ID estimator should be as insensitive to noise as possible. In order to test the robustness of M_2, the noisy Swiss rolls, presented in Subsection 4.1, were used and the results are displayed in Figure 5. On the left, a comparison between M_2, $\hat{m}_{LB}$ and $\hat{m}_{MG}$ was conducted to show how the mean and the standard deviation of the estimates (computed over 100 sets) change as the noise increases. Leaving aside the initial difference, the sensitivity of the three estimators appears to be similar. The right hand-side of the figure highlights that the number of points N has a low influence on the responsiveness of M_2 to noise. These results demonstrate that the behaviour of M_2 in presence of noise is not better or worse than that of the other two estimators.

4.5. Comparison Between the Estimators Using Real Data

M_2, D_2, $\hat{m}_{LB}$, $\hat{m}_{MG}$ and d_{cor} were used to estimate the ID of four real datasets from the UCI machine learning repository: (1) Housing values in suburbs of Boston, (2) Statlog (vehicle silhouettes), (3) Statlog (image segmentation) and (4) Combined cycle power plant. The output variables and

the duplicate events were removed. The number of points N and the number of variables E of the resulting datasets are indicated in Table 5. Notice that the true ID of real data is unknown and the goal of this subsection is to show that the results provided by M_2 are coherent with the theory and with the observations made for the synthetic data.

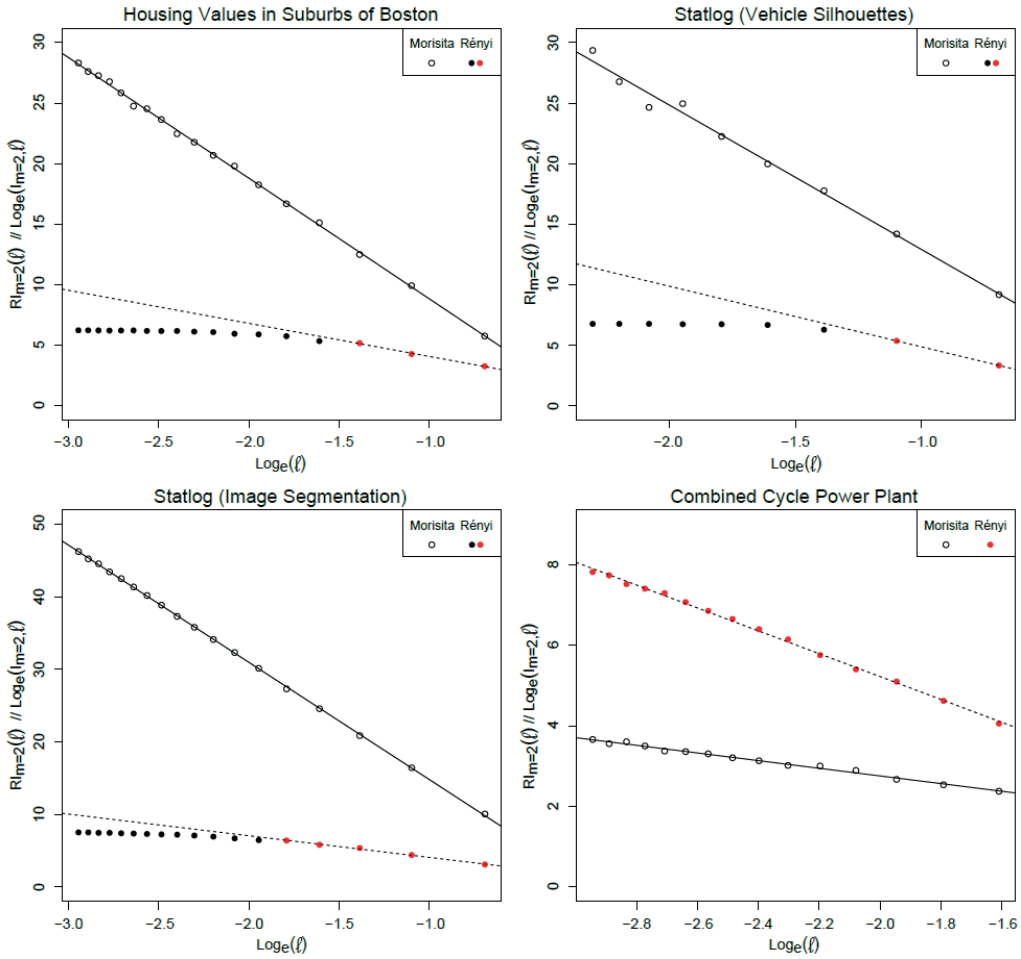

Figure 6: Application of $I_{2,\ell}$ and $RI_2(\ell)$ to four real datasets

Figure 6 shows the linear regressions required for the computation of M_2 and D_2. Here as well, $RI_2(\ell)$ appears to converge to a constant value as ℓ decreases. The departure from the power law denoted by the dashed line is

Real Data	Housing Values	Vehicle Silhouettes	Image Segmentation	Power Plant
N	506	846	2086	9527
E	13	18	19	4
M_2	3.02 $\ell^{-1} \in \{2, 3 \cdots, 19\}$	6.04 $\ell^{-1} \in \{2, 3 \cdots, 10\}$	2.84 $\ell^{-1} \in \{2, 3 \cdots, 19\}$	3.05 $\ell^{-1} \in \{5, 6 \cdots, 19\}$
D_2	2.73 $\ell^{-1} \in \{2, 3, 4\}$	5.02 $\ell^{-1} \in \{2, 3\}$	2.99 $\ell^{-1} \in \{2, 3, 4, 5, 6\}$	2.84 $\ell^{-1} \in \{5, 6 \cdots, 19\}$
$\hat{m}_{LB}$	3.64 $k \in \{4, 5, \cdots, 20\}$	5.63 $k \in \{3, 4, \cdots, 11\}$	3.54 $k \in \{10, 11, \cdots, 20\}$	3.11 $k \in \{4, 5, \cdots, 20\}$
$\hat{m}_{MG}$	2.98 $k \in \{4, 5, \cdots, 20\}$	5.57 $k \in \{3, 4, \cdots, 11\}$	2.98 $k \in \{4, 5, \cdots, 20\}$	2.78 $k \in \{4, 5, \cdots, 20\}$
d_{cor}	2.87 $0.05 \leq \delta \leq 0.5$	4.98 $0.1 \leq \delta \leq 0.4$	2.69 $0.04 \leq \delta \leq 0.15$	2.97 $0.02 \leq \delta \leq 0.16$

Table 5: Application of the ID estimators to four real datasets from the UCI machine learning repository

more pronounced when N is relatively low and the assumed ID rather high (according to the ID estimates of Table 5). This can be observed with the vehicle silhouette data. In contrast, the power plant dataset is characterized by a great number of points N and a relatively low ID (according to the final ID estimates of Table 5). As a result, the entire range of the parameter ℓ can be used to perform linear regression. But, regarding M_2, in each of the case studies, the solid line is a good approximation of the behaviour of $\log(I_{2,\ell})$ and all the values of ℓ can be systematically retained for ID estimation. This primacy of M_2 over D_2 is in accordance with the results found for the synthetic data, for which the Morisita estimator always provided equivalent or better ID estimates than D_q. The variability of $\log(I_{2,\ell})$ for small values of ℓ, observed in the case of the vehicle silhouettes, was also shown to have only a limited effect on the accuracy of M_2 (see Figures 3 and 4). Based on this examination of Figure 6, M_2 appears to adequately capture the scaling properties of the four real world datasets.

The final ID estimates of the real world datasets are given in Table 5. Although the results provided by D_2 were computed by using only the red points, they are less than those of M_2, except for the image segmentation data. It is also worth mentioning that D_2 yields 3.09 for the power plant dataset if it is calculated with only the three points on the right side of the plot, which is close to the estimate of M_2. These observations are coherent with the conclusions drawn from the synthetic data: for the lowest considered values of ℓ, the behaviour of $RI_2(\ell)$ can be evaluated by linear regression and the corresponding ID estimates tend to slightly underestimate both the ground truth and the values yielded by M_2 over the entire

range of examined scales. Besides, the difference between the two estimators fades away for rather large N and low theoretical ID (see Figures 3 and 4). The power plant and the vehicle silhouette datasets are real world illustrations of this last point. The number N of the former is large enough to fully characterize a 3-dimensional space (the assumed ID is about 3 and $N = 9527$), while the space where the latter resides is sparsely filled with points (the assumed ID is greater than 5 and $N = 846$). As a result, the range of ℓ and the final ID estimate of the two estimators differ more for the vehicle silhouette dataset than for the power plant one. In each plot of Figure 6, the difference is highlighted by the gap between the dashed line and the actual values of $RI_2(\ell)$ at the lowest scales. It is large in the case of the vehicle silhouettes, while it is non-existent for the power plant data. The two remaining datasets (i.e. the housing values and the image segmentation) are in an intermediate situation, as it is confirmed by the values given in Table 5. Thus, the comparison between M_2 and D_2 using real world data shows an excellent consistency with what was expected from the theory and the synthetic data analysis.

Regarding the distance-based methods, the estimates of df_{cor} are, as expected, lower than those provided by M_2. In contrast, the results of $\hat{m}_{LB}$ and $\hat{m}_{MG}$ tend to be less similar than for the synthetic data and they do not systematically underestimate the values of M_2. These slight differences, related to the complex specificities of the analysed datasets, do not question the good performance of the suggested estimator. Indeed, in each of the case studies, the values of M_2 are close to, at least, one of the two MLE-based results. And, finally, by comparison with the range of values yielded by all the other estimators, it can be definitely maintained that M_2 provides sensible ID estimates.

5. Conclusion

The Morisita estimator, M_m, is a new tool for estimating the Intrinsic Dimension (ID) of data. It is related to Rényi's generalized dimensions, D_q, for $m = q \geq 2$. M_m tended to provide better results than D_q on the synthetic data used in this study. This turned out to be particularly true for order 2 (i.e. $m = q = 2$) when the data points were sparsely distributed. The application to four real datasets from the UCI machine learning repository confirmed the good properties of the suggested estimator. From the perspective of pattern recognition, M_2 might be of great interest, since it

could be a good replacement for D_2 in algorithms, such as the fractal dimension algorithm [9, 12]. It might also open a new door to fractal supervised feature selection [13] of large datasets (our current work in progress), since its accuracy is coupled with a high computational efficiency.

M_2 was also compared with three distance-based estimators, namely d_{cor}, $\hat{m}_{LB}$ and $\hat{m}_{MG}$. It yielded good results when applied to the synthetic data and, by comparison, the real world applications revealed that it was able to provide reasonable ID estimates in challenging case studies. Therefore, in addition to the above-mentioned application in feature selection, the Morisita estimator of ID can be considered as a new tool for conducting advanced data mining tasks (e.g. dimensionality reduction [3], monitoring network analysis [36]) in many varied fields where the other estimators are commonly used (e.g. pattern recognition [27], physics [46, 39], cosmology [50], climatology [52] and ecology [47]).

Finally, it is also worth mentioning that the multipoint Morisita index is a ratio of probabilities deep-rooted in the field of spatial clustering analysis. Consequently, M_m can be viewed from a dual perspective, a fractal one and a statistical one, which helps to interpret the results.

6. Acknowledgements

The authors are grateful to the anonymous reviewers for their helpful and constructive comments that contributed to improving the paper. They also would like to thank Michael Leuenberger and Zhivko Taushanov for many fruitful discussions about machine learning and statistics.

References

[1] D. W. Scott, J. R. Thompson, Probability density estimation in higher dimensions, in: J. R. Gentle (Ed.), Proceedings of the Fifteenth Symposium on the Interface, Elsevier Science Publishers, North-Holland, 1983, pp. 173–179.

[2] R. Bellman, Adaptive Control Processes: A Guided Tour, Princeton University Press, Princeton (NJ), 1961.

[3] J. A. Lee, M. Verleysen, Nonlinear Dimensionality Reduction, Springer, New-York, 2007.

[4] C. J. C. Burges, Dimension reduction: A guided tour, Foundations and Trends in Machine Learning 2 (4) (2009) 275–365.

[5] J. B. Tenenbaum, V. de Silva, J. C. Langford, A global geometric framework for nonlinear dimensionality reduction, Science 290 (5500) (2000) 2319–2323.

[6] A. Lendasse, J. A. Lee, V. Wertz, M. Verleysen, Time series forecasting using CCA and Kohonen maps - application to electricity consumption, in: M. Verleysen (Ed.), Proceedings of ESANN 2000, 8th European Symposium on Artificial Neural Networks, Bruges, 2000, pp. 329–334.
[7] M. Kanevski, A. Pozdnoukhov, V. Timonin, Machine Learning for Spatial Environmental Data: Theory, Applications and Software, EPFL Press, Lausanne, 2009.
[8] S. Kaski, J. Peltonen, Dimensionality reduction for data visualization, Signal Processing Magazine 28 (2) (2011) 100–104.
[9] C. Traina Jr., A. J. M. Traina, L. Wu, C. Faloutsos, Fast feature selection using fractal dimension, in: Proceedings of the XV Brazilian Symposium on Databases (SBBD), 2000, p. 158–171.
[10] H. Zhang, C. Perng, Q. Cai, An improved algorithm for feature selection using fractal dimension, in: Proceedings of the Second International Workshop on Databases, Documents, and Information Fusion, 2002.
[11] E. P. M. De Sousa, C. Traina Jr., A. J. M. Traina, L. Wu, C. Faloutsos, A fast and effective method to find correlations among attributes in databases, Data Mining and Knowledge Discovery 14 (2007) 367–407.
[12] C. Traina Jr., A. J. M. Traina, C. Faloutsos, Fast feature selection using fractal dimension - Ten years later, Journal of Information and Data Management 1 (1) (2010) 17–20.
[13] D. Mo, S. H. Huang, Fractal-based intrinsic dimension estimation and its application in dimensionality reduction, IEEE Transactions on Knowledge and Data Engineering 24 (1) (2012) 59–71.
[14] J. G. Dy, C. E. Brodley, Feature Selection for Unsupervised Learning, Journal of Machine Learning Research 5 (2004) 845–889.
[15] J. G. Dy, Unsupervised feature selection, in: H. Liu, H. Motoda (Eds.), Computational Methods of Feature Selection, Chapman Hall/CRC, London/Boca Raton, 2013, pp. 29–55.
[16] S. Alelyani, J. Tang, H. Liu, Feature Selection for Clustering: A Review, in: C. C. Aggarwal, C. K. Reddy (Eds.), Data Clustering: Algorithms and Applications, Chapman Hall/CRC, London/Boca Raton, 2013, pp. 29–55.
[17] D. W. Aha, R. L. Bankert, A comparative evaluation of sequential feature selection algorithms, in: D. Fisher, H. J. Lenz (Eds.), Learning from Data: AI and Statistics, Springer, New-York, 1996, pp. 199–206.
[18] A. K. Jain, D. Zongker, Feature selection: Evaluation, application and small sample performance, IEEE Transactions on Pattern Analysis and Machine Intelligence 19 (2) (1997) 153–158.
[19] A. Blum, P. Langley, Selection of relevant features and examples in machine learning, Artificial Intelligence 97 (1-2) (1997) 245–271.
[20] L. C. Molina, L. Belanche, A. Nebot, Feature selection algorithms: A survey and experimental evaluation, in: Proceedings of 2002 IEEE International Conference on Data Mining (ICDM'02), Japan, 2002, pp. 306–313.
[21] I. Guyon, A. Elisseeff, An introduction to variable and feature selection, Journal of Machine Learning Research 3 (2003) 1157–1182.
[22] I. Guyon, S. Gunn, M. Nikravesh, L. A. Zadeh, Feature extraction: Foundations and Applications, Springer, Berlin, 2006.

[23] T. Marill, D. M. Green, On the effectiveness of receptors in recognition systems, IEEE Transactions on Information Theory 9 (1) (1963) 11–17.
[24] P. Pudil, J. Novovičová, J. Kittler, Floating search methods in feature selection, Pattern Recognition Letters 15 (11) (1963) 1119–1125.
[25] S. M. Vieira, M. C. Sousa, T. A. Runkler, Ant colony optimization applied to feature selection in fuzzy classifiers, Lecture Notes in Computer Science 4529 (2007) 778–788.
[26] X. Wang, J. Yang, X. Teng, W. Xia, R. Jensen, Feature selection based on rough sets and particle swarm optimization, Pattern Recognition Letters 28 (4) (2007) 459–471.
[27] F. Camastra, Data dimensionality estimation methods: a survey, Pattern Recognition 36 (12) (2003) 2945 – 2954.
[28] E. Levina, P. J. Bickel, Maximum likelihood estimation of intrinsic dimension 17, in: Advances in Neural Information Processing Systems, Vol. 17, The MIT Press, Cambridge (USA), 2004.
[29] J. Golay, M. Kanevski, C. D. Vega Orozco, M. Leuenberger, The multipoint Morisita index for the analysis of spatial patterns, Physica A 406 (2014) 191–202.
[30] S. H. Hurlbert, Spatial Distribution of the Montane Unicorn, Oikos 58 (3) (1990) 257–271.
[31] M. Morisita, Measuring of the Dispersion of Individuals and Analysis of the Distributional Patterns, Memoires of the Faculty of Science (Serie E), Kyushu University 2 (4) (1959) 215–235.
[32] A. K. Jain, R. C. Dubes, Algorithms for Clustering Data, Prentice-Hall, New-Jersey, 1988.
[33] J. Theiler, Estimating fractal dimension, Journal of the Optical Society of America 7 (6) (1990) 1055 – 1073.
[34] B. Kégl, Intrinsic dimension estimation using packing numbers, in: Advances in Neural Information Processing Systems, Vol. 14, The MIT Press, Cambridge (USA), 2002.
[35] B. B. Mandelbrot, The Fractal Geometry of Nature, W.H. Freeman, San Francisco, 1983.
[36] S. Lovejoy, D. Schertzer, P. Ladoy, Fractal Characterization of Inhomogeneous Geophysical Measuring Networks, Nature 319 (6048) (1986) 43–44.
[37] T. G. Smith, W. B. Marks, G. D. Lange, W. H. Sheriff, E. A. Neale, A Fractal Analysis of Cell Images, Journal of Neuroscience Methods 27 (2) (1989) 173–180.
[38] K. Falconer, Fractal Geometry: Mathematical Foundations and Applications, 2nd Edition, Wiley, Chichester (UK), 2003.
[39] H. G. E. Hentschel, I. Procaccia, The Infinite Number of Generalized Dimensions of Fractals and Strange Attractors, Physica D 8 (3) (1983) 435–444.
[40] P. Grassberger, Generalized Dimensions of Strange Attractors, Physics Letters A 97 (6) (1983) 227–230.
[41] G. Paladin, A. Vulpiani, Anomalous Scaling Laws in Multifractal Objects, Physics Reports 156 (4) (1987) 147–225.
[42] T. Tel, A. Fülöp, T. Vicsek, Determination of Fractal Dimensions for Geometrical Multifractals, Physica A 159 (1989) 155–166.
[43] T. Vicsek, Fractal Growth Phenomena, World Scientific, Singapore, 1993.

[44] A. Rényi, Probability Theory, Akadémiai Kiadò, Budapest, 1970.
[45] J. Balatoni, A. Rényi, On the Notion of Entropy, Publ. Math. Inst. Hungarian Acad. Sci (1) (1956) 5–40, english translation in *Selected Papers of A. Rényi*, Budapest, vol. 1 (1976), 558.
[46] P. Grassberger, I. Procaccia, Characterization of strange attractors, Physical Review Letters 50 (5) (1983) 346–349.
[47] L. Seuront, Fractals and Multifractals in Ecology and Aquatic Science, CRC Press, Boca Raton (USA), 2010.
[48] Y. Chen, Multifractals of central place systems: Models, dimension spectrums, and empirical analysis, Physica A 402 (2014) 266–282.
[49] M. M. Dubovikov, N. V. Starchenko, M. S. Dubovikov, Dimension of the minimal cover and fractal analysis of time series, Physica A 339 (3-4) (2004) 591–608.
[50] S. Borgani, G. Murante, A. Provenzale, R. Valdarnini, Multifractal Analysis of the Galaxy Distribution: Reliability of Results from Finite Data Sets, Physical Review E 47 (6) (1993) 3879–3888.
[51] S. Lovejoy, D. Schertzer, A. Tsonis, Functional Box-counting and Multiple Elliptical Dimensions in Rain, Science 235 (4792) (1987) 1036–1038.
[52] S. Lovejoy, D. Schertzer, The Weather and Climate, Cambridge University Press, New-York, 2013.
[53] Q. Huang, J. R. Lorch, R. C. Dubes, Can the fractal dimension of images be measured?, Pattern Recognition 27 (3) (1994) 339–349.
[54] J. B. Florindo, O. M. Bruno, Fractal descriptors based on the probability dimension: A texture analysis and classification approach, Pattern Recognition Letters 42 (2014) 107–114.
[55] F. Camastra, A. Vinciarelli, Estimating the intrinsic dimension of data with a fractal-based method, IEEE Transactions on Pattern Analysis and Machine Intelligence 24 (10) (2002) 1404–1407.
[56] J. Eberhardt, Estimating Intrinsic Dimension, University of Minnesota, Duluth, 2007.
[57] A. Asensio Ramos, H. Socas-Navarro, A. López Ariste, M. J. Martínez González, The intrinsic dimensionality of spectropolarimetric data, The Astrophysical Journal 660 (2007) 1690–1699.
[58] J. C. MacKay, D, Z. Ghahramani, Comments on "maximum likelihood estimation of intrinsic dimension" by E. Levina and P. Bickel @ONLINE (Jan. 2005).
URL `http://http://www.inference.phy.cam.ac.uk/mackay/dimension/`
[59] D. Tuia, M. Kanevski, Envrionmental Monitoring Network Charaterization and Clustering, in: M. Kanevski (Ed.), Advanced Mapping of Environmental Data: Geostatistics, Machine Learning and Bayesian Maximum Entropy, Iste/Wiley, London/Hoboken(USA), 2008, pp. 19–46.
[60] M. Kanevski, M. Maignan, Analysis and Modelling of Spatial Environmental Data, EPFL Press, Lausanne, 2004.
[61] Y. Shi, C. Gong, Multifractality of a cartesian product of two fractals, Communications in Theoretical Physics 23 (1995) 245–248.
[62] D. B. Ripley, Spatial Statistics, Wiley, New-York, 1981.